Keeping Up with the Joneses: Other-regarding Preferences and Endogenous Growth

Luke Petach, Daniele Tavani*

Colorado State University

September 29, 2017

Abstract

We study a series of sustained growth models in which households' preferences are affected by the consumption of other households as summarized by average consumption. In endogenous growth models, the equilibrium paths involve lower savings and lower growth than the corresponding efficient paths. Both savings and growth are inversely related to the extent of social preferences. In semi-endogenous models, other-regarding preferences have no growth effects, but have positive level effects on the long-run research intensity, because they increase the market size for new potential monopolists in the intermediate goods sector. We then use Consumer Expenditure Survey data to test the extent to which consumption is other-regarding: our identification strategy relies on a two-stage estimator that uses the Tax Reform Act of 1986 and the Omnibus Budget Reconciliation Act of 1993 as a positive and a negative consumption shocks to top incomes respectively. In the first stage, we use a difference-in-difference approach to exploit the exogenous variation in consumption caused by federal tax reform. We then use the predicted values for average within-cohort consumption by income deciles as an instrumental variable to estimate the extent of social preferences. Our results point toward significant long-run 'keeping up' effects on the order of 30%.

*Preliminary draft. Corresponding author. Department of Economics, Colorado State University. 1771 Campus Delivery, Fort Collins, CO 80523-1771. Email: Daniele.Tavani@Colostate.edu.

1. Introduction

In modeling consumption and saving decisions, the analysis of long-run growth has typically focused on atomistic, forward-looking households making choices in isolation, and whose preferences depend only on the individual fruition of goods and services. However, economists have long suspected that individual choices may be affected by the comparison with others. Inter-household comparisons are the core of Thorstein Veblen's analysis of conspicuous consumption (Veblen, 1899, 2007); James Duesenberry, on the other hand, has highlighted both the importance of consumption and saving decisions made by others —the *relative income hypothesis*— and the role played by habit persistence in consumption patterns (Duesenberry, 1949). In an empirical test that distinguishes between the two motives, Ravina (2007) uses US microeconomic data to find strong evidence —on the order of 30%— of the role played by 'external habits' as well as 'internal habits' on individual household consumption. Relative income plays an important role on household self-reported happiness, as documented by Dynan and Ravina (2007).

One common element in the literature on consumption decisions under other-regarding preferences (Carroll et al., 1997; Ljunqvist and Uhlig, 2000; Turnovsky *et al.*, 2004; Liu and Turnovsky, 2005; Ravina, 2007; Dynan and Ravina, 2007; Alvarez-Cuadrado and Van Long, 2011) is that the postulated effect of peer choices on individual behavior is a 'deep' feature of the economic environment: assumptions on preferences, together with technology and endowments, are the building blocks of any dynamic general equilibrium framework. Thus, other-regarding preferences should be regarded as a long-run characteristic of the economy. And yet, long-run growth is unaffected by 'keeping up' effects in the Neoclassical growth model (NGM): because of diminishing returns to capital per worker, the economy reaches a fully exogenous balanced growth path (BGP) in which consumption and capital stock —both measured in effective worker units— are constant over time, and GDP per capita grows at the exogenous growth rate of labor productivity (Turnovsky *et al.*, 2004). Other-regarding households typically consume more and save less than isolated households, but saving behavior only has level effects on long-run income and consumption per effective worker. Thus, while 'keeping up' effects usually matter along the transitional dynamics as they slow down consumption growth over time (Turnovsky *et al.*, 2004), they become irrelevant as soon as the economy approaches balanced growth. Furthermore, the presence of other-regarding preferences does not affect the welfare properties of the BGP, which is Pareto-efficient exactly like in the isolated household case. The intuition for the irrelevance of peer effects in Neoclassical growth is very simple: in the long run, the theory assigns no role for investment (and thus saving) decisions in shaping the growth rate of the economy along a BGP.

Conversely, saving behavior plays a prominent role in endogenous growth models based on increasing returns: a reduction in the determinants of savings translates into a permanently lower growth rate. (Romer, 1986; Lucas, 1988; Barro, 1990; Romer, 1990; Grossman and Helpman, 1991). The empirical evidence provided by Bernanke and Gürkaynak (2001) has shown a strong correlation between long-run growth and behavioral parameters such as the saving rate. This result is taken as lending support for endogenous, as opposed to exogenous, growth models. If one buys the endogeneity of the growth process, and 'keeping up' behavior decreases the saving rate of the economy, then it is bound to have permanent long-run growth effects. However, such conclusion depends on the postulated form of the dependence of individual utility on the reference variable: in the 'outward-looking' case analyzed by Carroll et al. (1997) —the closest paper to the present analysis— households care about the ratio of their consumption to the reference group. They found that a shock to capital decreases saving and growth, but only temporarily.

More recently,Ljunqvist and Uhlig (2000) and Dynan and Ravina (2007) have posited that the relevant comparison is not the ratio but the *difference* with consumption of the reference group, the latter multiplied by a 'social preference' parameter $\theta \in (0, 1)$. This type of preferences captures the notion of 'keeping up with the Joneses', or external habits, and it forms the starting point of our analysis. As a benchmark, we consider a transparent AK growth model with 'keeping up' preferences. In balanced growth, the policy function that traces the utility-maximizing consumption path will not only depend on the household's asset base as in the standard version of the model (Acemoglu, 2009, Chapter 11), but will also have the features of a best-response function to average consumption. In equilibrium, however, the two have to be equal, similarly to the notion of a Nash equilibrium. As a consequence, the equilibrium growth rate is inversely related to the extent of social preferences: the stronger the 'keeping up' effect, the lower the saving rate, the lower accumulation and growth. The resulting aggregate consumption rule will not only feature household wealth as its argument, but also the household's current disposable income, with the 'marginal propensity to consume' out of income being proportional to the extent of social preferences.

It is then straightforward to show that the equilibrium growth path of this model is not Pareto-efficient. A benevolent planner would internalize that the representative household's consumption and average consumption have to be equal to each other, thus annihilating the perverse effect of social preferences on capital accumulation. Decentralizing the growth path can in principle be achieved by levying taxes on consumption and subsidizing household incomes in order to offset the 'keeping up' effect. The size of the resulting income subsidy rises in the extent to which preferences are other-regarding.

Because the basic tenet of the importance of savings for growth is common across a variety of endogenous growth models, it makes sense to evaluate the role of consumption peer effects in other sustained growth frameworks. Growth models based on increasing returns —be those driven by productive public services (Barro, 1990), human capital accumulation (Lucas, 1988) or product-variety (Romer, 1990)— are characterized by inefficiently low growth rates along their respective balanced growth paths. In all of these models, inefficiencies arise because there are spillovers that are not internalized by the individual households. We show that other-regarding preferences produce an additional source of inefficiency that lowers the equilibrium growth rate relative to the original models without consumption peer effects.

Finally, we consider the implications of other-regarding behavior for semi-endogenous growth à la Jones (1995), where the long-run growth rate is determined within the model, but is only a function of policy-invariant parameters such as population growth and the technology governing the production of new ideas. Here, unsurprisingly, 'keeping up' behavior has no effect on long-run growth. However, and differently from the NGM, consumption peer effects do produce level effects on the equilibrium allocation of workers between final goods production and R&D. The more preferences are other-regarding, the more researchers relative to manufacturing workers will be hired in equilibrium. The reason is that peer effects generate an upward shift in the demand for the final good that increases the size of the market for the producers of new varieties of idea-based goods, thus providing an incentive to allocate more resources into R&D activities.

In order to empirically validate our hypothesis, we use Consumer Expenditure (CEX) data to test the response of individual consumption to average consumption. We follow Agarwal and Qian (2014) and Misra and Surico (2014) in using fiscal policy shocks to identify variation in consumption. Our identification strategy exploits the Tax Reform Act of 1986 (TRA86) and the Omnibus Budget Reconciliation Act of 1993 (OBRA93). TRA86 was characterized by: (i) tax cuts to income-earners facing top marginal rates;[1] (ii) increased tax incentives for owner-occupied housing by expanding the mortgage interest-rate tax deduction, and (iii) a reduction in the total number of tax brackets. All these reforms positively effect household disposable incomes and provide an impetus for consumption. In contrast, OBRA93 raised marginal rates for top-income earners, reversing some of the effects of TRA86. In both cases, we use a two-stage process in order to address the obvious endogeneity problem that affects any measure of average consumption as the main explanatory variable in a regression equation. After dividing the sample into income deciles, we use a difference-in-difference estimator that exploits the exogenous variation in consumption caused by federal tax reform. Using these estimation results, we calculate predicted average within-cohort consumption which is then used as instrument to obtain an estimate of the extent of social preferences in a second-stage regression.

The use of both a positive and negative shock allows us to test multiple hypotheses regarding the social preference parameter. The first question is whether the available evidence points toward social preferences to be persistent or not. If the estimated consumption peer effects after the two policy shocks are close to each other, we can reasonably conclude that social preferences are relevant across different instances of policy shocks, rather than being merely a transient phenomenon. This is exactly the case in our analysis: the point estimates for TRA86 and OBRA93 are around .33 and .28 respectively, highly statistically significant, and we fail to reject the null of the two values being equal to each other. We conclude that consumption peer effects can be treated as a long-run feature of the economy. Using an instrument for current average consumption also marks a difference between our contribution and the empirical strategy used Turnovsky *et al.* (2004), who specify the reference stock as a weighted average of historic economy-wide average consumption in order to test the 'catching up with the Joneses' hypothesis by Ljunqvist and Uhlig (2000). The intuition behind the notion of 'keeping up' suggests that individual consumption should depend on

[1] A history of top marginal rates is available from the Internal Revenue Service: https://www.irs.gov/uac/soi-tax-stats-historical-data-tables.

either the current consumption or income of the reference group.

We confirm this result by running a number of robustness checks including placebo tests for both TRA86 and OBRA93, assessing the strength of the social preference parameter by income decile via quantile regression, and estimating restricted sample specifications to control for sub-groups that may be biasing the regression coefficient of interest. These checks confirm the strength of our initial result. Additionally, the results from quantile regression suggest that the strength of social preferences in consumption decline as one moves up the income distribution, indicating that the existence of such preferences may have important implications for inequality over time. This result puts our analysis in line with the findings of Alvarez-Cuadrado and Van Long (2011), whose theoretical framework delivers higher saving rates for upper deciles of the income distribution.

2. A Benchmark, Log-Utility AK Model

We consider a one-good closed economy in continuous time. Output per worker y is produced using the simplest technology that generates sustained growth: $y = Ak$. In order to facilitate the comparison with idea-based endogenous growth models, we assume that capital stock does not depreciate and that population is constant.

Suppose in standard fashion that there is a representative household choosing a sequence of consumption streams $\{c(t)\}_{t\in[s,\infty)}$ so as to maximize the present-discounted value of its lifetime utility. We denote the rate of time preference, constant throughout, by $\rho > 0$. Further, and in order to focus on growing economies, we assume that $A > \rho$.

Following Dynan and Ravina (2007) and Alvarez-Cuadrado and Van Long (2011), we assume that households' choices about consumption are affected by the consumption choices made by other households, as summarized by average consumption $\bar{c}(t)$. Specifically, we assume that the argument of the utility function is the difference between consumption and average consumption, the latter multiplied by a 'social preference' parameter $\theta \in [0, 1)$. The magnitude of the social preference parameter determines how much the household attempts to 'keep up with the Joneses:' $\theta = 0$ implies that the household is able to isolate itself from peer effects, while increasing values of θ imply more other-regarding behavior. Average consumption is taken as a given by each household in its utility-maximization program. As such, it has the features of an externality: the household does not take into account the fact that its consumption decisions also affect average consumption in the economy.

The simplest possible utility function for the problem at hand is logarithmic in its argument. It is a textbook result that, in a 'canonical' AK model with log utility and isolated households, balanced growth requires households to consume a constant fraction of their wealth, the fraction being equal to the pure rate of time preference (Acemoglu, 2009, Chapter 11). To make our analysis as close as possible to this case, we assume that instantaneous preferences are given by:

$$u(c,\bar{c}) = \ln[c-\theta\bar{c}], \theta \in (0,1). \tag{1}$$

For the sake of simplicity, we rule out the possibility that the household runs into debt in order to finance its consumption, so that its budget constraint in per-capita terms is $\dot{k}(t) = Ak(t) - c(t)$. A standard optimal control exercise leads to the following dynamical system in consumption and capital stock, both in per-capita terms:

$$\frac{\dot{c}(t)}{c(t)} = \frac{c(t)-\theta\bar{c}(t)}{c(t)}\left[A-\rho\right] \tag{2}$$

$$\frac{\dot{k}(t)}{k(t)} = A - \frac{c(t)}{k(t)} \tag{3}$$

2.1. Balanced Growth

A balanced growth path (BGP) for the economy is defined as sequences of per-capita consumption and capital stock such that $\dot{c}/c = \dot{k}/k$ at all times. Along a BGP, the household's choice of consumption has the features of a best-response function to average consumption. Setting the right hand side of (2) and (3) equal to each other, and solving for consumption per capita, we find:[2]

$$c[\bar{c}(t), k(t)] = \frac{\rho k(t)}{2} + \frac{\left\{[\rho k(t)]^2 + 4\theta\bar{c}(t)\left[A-\rho\right]k(t)\right\}^{1/2}}{2}. \tag{BR}$$

[2]We restrict our attention to the solution that returns non-negative consumption for any non-negative value of average consumption. It is easy to check that the other solution of the quadratic equation leading to (BR), in fact, has zero intercept and negative consumption for any level of $\bar{c}$.

The choice of consumption has intercept equal to $\rho k(t)$ and is increasing and concave in average consumption. It is also apparent that the presence of social preferences makes the representative household deviate from the standard BGP policy function. In fact, under $\theta = 0$, equation (BR) reduces to its textbook counterpart:

$$c^*[k(t)] = \rho k(t). \tag{4}$$

This no-externality case constitutes a benchmark of Pareto-efficiency for our analysis, as it is shown below.

2.2. Equilibrium

An equilibrium path for the economy is a BGP such that $c(t) = \bar{c}(t)$ for all t. This definition is close to the notion of a Nash equilibrium. At an equilibrium path, household's consumption fulfills:

$$c^E[k(t); \theta] = (1-\theta)\rho k(t) + \theta A k(t) \tag{5}$$

Notice that, as long as $\theta \in (0, 1)$, aggregate consumption is sensitive to current income Ak. Equation (5) actually looks like an undergraduate textbook Keynesian consumption function, with 'autonomous consumption' given by $(1-\theta)\rho k$ and the 'marginal propensity to consume' equal to the extent of social preferences. The equilibrium growth rate is

$$g = (1-\theta)[A - \rho]. \tag{G}$$

Finally, the equilibrium saving rate of the economy is constant, and inversely related to the extent of social preferences:

$$s = \frac{\dot{k}}{\dot{k} + c} = (1-\theta)\left[1 - \frac{\rho}{A}\right]. \tag{S}$$

2.3. Welfare Analysis

The equilibrium growth path of this economy is not Pareto-efficient. To see this, consider the problem faced by a benevolent planner solving the household's utility maximization problem under the additional constraint that $c(t) = \bar{c}(t)\ \ \forall t$. The planner, thus, internalizes that the utility function boils down to $u[c(t)] = \ln(1-\theta) + \ln c(t)$. The optimal control solution for this problem satisfies equation (3) and the consumption Euler equation:

$$g^* = \frac{\dot{c}}{c} = A - \rho. \tag{G*}$$

It is obvious that $g^* > g \ \forall\theta \in (0,1)$. Further, the policy function on the efficient path is (4). Finally, it is straightforward to show that $c^*(t) > c(t) \ \forall\theta \in (0,1)$, and that the efficient saving rate is:

$$s^* = 1 - \frac{\rho}{A} > s \, \forall\theta \in (0,1). \tag{S*}$$

Figure 1 illustrates the equilibrium level of consumption in comparison with the efficient level.

[Figure 1 about here.]

2.4. Decentralization

How can the Pareto-efficient path be decentralized as an equilibrium path? A public authority could tax consumption and rebate tax revenues in the form of an income subsidy to households, in order to offset the effect of social preferences. Suppose that the government chooses a proportional income subsidy, $\sigma y, \sigma \in (0,1)$, financed by a consumption tax $\tau c, \tau \in (0,1)$. At a balanced budget, $\sigma y = \tau c$. The representative household's budget constraint modifies to

$$Ak(t)(1+\sigma) = c(t)(1+\tau) + \dot{k}(t). \tag{6}$$

Solving for the consumption Euler equation gives

$$\frac{\dot{c}}{c} = \left(\frac{c - \theta\bar{c}}{c}\right)[A(1+\sigma) - \rho] \tag{7}$$

while, because of the balanced budget requirement for the government, the aggregate accumulation equation is still given by (3). Again, we look at an equilibrium for the economy where $c(t) = \bar{c}(t)$ at all times. The growth rate becomes simply

$$\frac{\dot{c}}{c} = (1-\theta)[A(1+\sigma) - \rho],$$

and the value for the subsidy that decentralizes the efficient path is found as

$$\sigma = \frac{\theta}{1-\theta}\left(1 - \frac{\rho}{A}\right) = \frac{\theta}{1-\theta} s^* \tag{8}$$

The subsidy is increasing and convex in θ, the extent to which preferences are other-regarding.

2.4.1. CRRA Preferences

The properties of the equilibrium path *vs.* the efficient path carry over to the more general case of constant-relative-risk-aversion (CRRA) preferences. Omitting the time dependence for notational simplicity, and denoting the relative risk-aversion parameter by $\gamma > 0, \gamma \neq 1$, the consumption Euler equation becomes:

$$\frac{\dot{c}}{c} = \left(\frac{c - \theta\bar{c}}{\gamma c}\right)[A - \rho] \tag{9}$$

Appendix A shows that the best response function is again increasing and concave in average consumption, and that the equilibrium consumption and growth rate are respectively larger and smaller than their efficient counterparts.

2.5. Comparison with Neoclassical Growth

Suppose, as is the case in the basic NGM, that the intensive production function exhibits diminishing returns to capital stock per worker: $y = Ak^\alpha, \alpha \in (0, 1)$. With diminishing returns, other-regarding preferences are irrelevant in the long-run, but only affect the speed of convergence to the steady state of the model. The equilibrium Euler equation is

$$\frac{\dot{c}}{c} = (1-\theta)\left[\alpha A k^{\alpha-1} - \rho\right]$$

and yields the steady state value for capital stock per person $k_{ss} = [\alpha A/\rho]^{1/(1-\alpha)}$, independent of social preferences. The same is true for consumption per capita $c_{ss} = A^{\frac{1}{1-\alpha}}(\alpha/\rho)^{\frac{\alpha}{1-\alpha}}$, independent of peer-effects. In this case, then, the decentralized steady state allocation is Pareto-efficient. However, the speed of convergence to the steady state is negatively affected by the extent of social preferences, as shown in Appendix B.

3. Other Sustained Growth Frameworks

3.1. Productive Public Services

In a seminal paper, Barro (1990) has provided a justification for endogenous growth that hinges on the role of public services that directly affect the production possibilities of the economy. His production technology combines physical capital k and a flow of public goods π through the Cobb-Douglas specification: $y = Ak^{\alpha}\pi^{1-\alpha}$. The public good is financed through an income tax, and the government runs a balanced budget: $\pi = \tau y$, but is taken as given by individual firms when making production decisions. The main results of Barro's analysis are that: (i) the growth-maximizing tax rate is equal to the output-elasticity of the public good; (ii) the growth-maximizing tax rate is the same in the decentralized growth path and in the efficient growth path, but (iii) the equilibrium path involves a lower growth rate than the efficient path.

Embedding social preferences into this model gives the following Euler equation for the individual household:

$$\frac{\dot{c}}{c} = \frac{c - \theta\bar{c}}{\gamma c}\left[(1-\tau)A(\pi/k)^{1-\alpha} - \rho\right]$$

which, once the equilibrium condition $c = \bar{c}$ and the balanced-budget requirement for the government are imposed, gives the equilibrium growth rate as:

$$g = \frac{1-\theta}{\gamma}\left[\alpha(1-\tau)A^{1/\alpha}\tau^{\frac{1-\alpha}{\alpha}} - \rho\right] \tag{10}$$

The equilibrium growth rate in the Barro (1990) model g^B is easily found by setting $\theta = 0$ in (10). The efficient growth rate is, instead,

$$g^* = \frac{1}{\gamma}\left[(1-\tau)A^{1/\alpha}\tau^{\frac{1-\alpha}{\alpha}} - \rho\right] \tag{11}$$

and the growth paths can be ranked as follows: $g < g^B < g^*$. All three growth rates are maximized by levying the same tax rate, equal to the elasticity of output to the public good: $\tau^* = 1 - \alpha$. But social preferences produce an additional source of market failure that lowers the growth rate relative to the isolated household case studied in Barro (1990).

3.2. Human Capital

The same kind of conclusions holds in a model with human capital overcoming the extent of diminishing returns to physical capital, such as the one by Lucas (1988). Such model, in balanced growth, delivers a constant marginal product of aggregate capital and therefore sustained growth: the corresponding growth rate will be formally similar to g in equation (G), which is negatively affected by the extent of social preferences. Denoting the growth rate in the Lucas (1988) model by g^L, then $g < g^L$. It is then possible to rank growth paths, similarly to what above: $g < g^L \leq g^*$, with strict inequality if the average human capital causes positive externalities in production, as argued by Lucas (1988).

3.3. Product Variety

The endogenous growth literature has emphasized the role of ideas —both with regards to product quality (Grossman and Helpman, 1991; Aghion and Howitt, 1992) and product variety (Romer, 1990)— in the growth process. Due to the formal similarities between the models, we focus here only on product variety. In the simplest framework, there are no capital goods: instead, at time t there is a continuum of mass $A(t)$ of varieties of intermediate goods, $x(i,t)$ that are combined with labor to produce the final good. Furthermore, following Romer (1990), let us assume that the total labor force L, constant over time, can either be employed in the production of the final good Y or in the production of new ideas. the labor market clearing condition requires that $L = L_Y + L_A$. The production function for the final good modifies as follows:

$$Y(t) = L_Y^{1-\alpha} \int_0^{A(t)} x^{\alpha}(i,t)di \tag{12}$$

Each intermediate good is supplied to final goods producers by a monopolistically competitive firm. Firms in the intermediate goods sector face an inverse demand function requiring the price of good i to equal the marginal product of the corresponding variety of intermediate products in production of the final good: $p(x(i)) = \alpha L_Y^{1-\alpha} x(i)^{\alpha-1}$ at time t. One unit of intermediate good in sector i can be produced with a unit of final good with unit cost equal to 1. Thus, profit maximization leads to the following supply function and profit function, both equal across all intermediate monopolists, and both proportional to employment in the final-good sector:

$$x(i) = x = \alpha^{\frac{2}{1-\alpha}} L_Y \ \forall i; \Pi = \frac{1-\alpha}{\alpha} x \ \forall i. \tag{13}$$

Next, both types of workers earn the same real wage $w(t)$ at time t, in turn equal to the marginal product of labor in the production of the final good. Thus, using (13),

$$w(t) = (1-\alpha)A(t)\alpha^{\frac{2\alpha}{1-\alpha}} L_Y \tag{14}$$

The R&D business is competitive and characterized by free entry. The rate at which ideas grow over time is linear in the number of workers in the R&D sector:

$$g_A \equiv \frac{\dot{A}(t)}{A(t)} = \lambda L_A \tag{15}$$

Each potential monopolist can invest in developing a new variety of intermediate goods, or in a risk-free asset $a(t)$ with rate of return equal to $r(t)$ at time t. A standard arbitrage condition requires the value of the asset to be equal to the ratio of monopoly profits Π over the rate of return r. The zero-profit condition in R&D can therefore be written as $\frac{\Pi}{r(t)} \lambda A(t) L_A = w(t) L_A$ which, using (14) gives the (constant) rate of return as:

$$r = \lambda \alpha L_Y \tag{16}$$

The representative household maximizes the present-discounted value of the instantaneous utility streams under the budget constraint requiring consumption $c(t)$ plus accumulation $\dot{a}(t)$ not to exceed the sum of its wage income $w(t)$ and interest income $r(t)a(t)$, ruling out depreciation for simplicity. With constant intertemporal elasticity of substitution, the consumption Euler equation is:

$$g_c \equiv \frac{\dot{c}}{c} = \frac{c - \theta \bar{c}}{\gamma c} [r - \rho] \tag{17}$$

In balanced growth, it must be the case that $g_c = g_A = g$. Thus, the rate of return is $r = \alpha(\lambda L - g)$, and the growth rate is

$$g = \frac{c - \theta\bar{c}}{\gamma c + \alpha(c - \theta\bar{c})}[\alpha\lambda L - \rho]$$

Imposing the Nash equilibrium condition $c = \bar{c}$, we obtain

$$g = \frac{1 - \theta}{\gamma + \alpha(1 - \theta)}[\alpha\lambda L - \rho], \tag{18}$$

which makes it clear that social preferences matter in this case, too. The growth rate of the standard product-variety model without social preferences is easily found by setting $\theta = 0$. Call this growth rate g^R which, however, is not Pareto-efficient because R&D firms do not take into account the non-rival nature of ideas. The Pareto-efficient growth rate would be $g^* = \gamma^{-1}(\lambda L - \rho)$. Hence, the presence of social preferences once again produces an additional source of inefficiency that lowers the growth rate of the product variety model. In fact, we have the following unambiguous ranking of growth paths:

$$g < g^R < g^*.$$

3.4. Semi-Endogenous Growth

In order to remove the counterfactual implication of a scale effect from the size of the economy (L) to the growth rate, Jones (1995) introduced a less-than-linear spillover from past technology to new ideas. The production function of ideas can be modified as follows:

$$\dot{A}(t) = \lambda L_A A(t)^{\phi}, \phi \in (0, 1) \tag{19}$$

In balanced growth, and with population growing at a constant rate $n > 0$, this model delivers the semi-endogenous —that is, determined within the model but policy-invariant— growth rate $g = n/(1 - \phi)$. Because the growth rate does not depend on savings, the extent of social preferences has no effect on growth in this framework. However, it is easy to show that there is a level effect of social preferences on the amount of R&D performed in balanced growth. Define the ratio of researchers to final good-producing workers as

$L_Y/L_A \equiv \chi$. The zero-profit condition in R&D implies that

$$\chi = \frac{r}{\alpha}\frac{1}{g},$$

and therefore the Euler equation for consumption at an equilibrium where $c = \bar{c}$ can be written as

$$\frac{\dot{c}}{c} = g = \frac{1-\theta}{\gamma}[\alpha\chi g - \rho] \tag{20}$$

which, using the growth rate derived above, pins down the balanced-growth research intensity as:

$$\chi = \frac{1}{\alpha}\left[\frac{\gamma}{1-\theta} + \frac{\rho(1-\phi)}{n}\right] \tag{21}$$

The balanced-growth research intensity in the Jones (1995) model, call it χ^J, is simply found by setting $\theta = 0$ and, interestingly enough, it is *smaller* than χ. The intuition is straightforward but insightful: the extent of social preferences produces a shift in the demand for the final good, that can be met by an increase in the number of existing varieties of intermediate products. This effect increases the size of monopoly profits for new intermediate goods producers, which then have an incentive to allocate more resources to R&D activities.[3]

4. Estimation

In order to empirically validate the model presented above, a natural question to ask is to what extent do households actually take into account the consumption decisions of others when solving for the decentralized consumption path. Answering this question amounts to evaluating the magnitude of the social preference parameter θ. To motivate our estimation strategy, we consider a general form of (BR):

$$c = c[\overline{c}(t), k(t)] \tag{22}$$

[3]The same is true about the other parameters that increase current consumption at the expenses of future consumption, such as the intertemporal elasticity of substitution γ and the discount rate ρ.

Where (22) is increasing and concave in its arguments. Its first-derivative can be written:

$$\frac{\partial c}{\partial \overline{c}} = g(\theta, \overline{c}(t), k(t)) \tag{23}$$

and has the property $g(\theta, \overline{c}(t), k(t)) = 0$ if $\theta = 0$. Given the above features, we propose a regression of the following form:

$$c_{ijt} = \theta \overline{c}_{jt} + \gamma y_{it} + X_{it}^T \beta + \epsilon_{it} \tag{2S}$$

where observations are indexed by household, income cohort (defined by deciles of the income distribution), and time. c_{ijt}, $\overline{c}_{jt}$, y_{it}, ϵ_{it}, and X_{it} denote household consumption, average consumption of income cohort j, household income, an idiosyncratic error, and a vector of household-specific deterministic characteristics, respectively. In this case θ measures the marginal propensity to consume out of *average* cohort consumption.

Estimation of (2S) is immediately problematic due to the endogeneity arising from the obvious correlation between c_{ijt} and $\overline{c}_{jt}$. To solve the identification problem, we propose to exploit two instances of federal income tax reform as sources of plausibly exogenous variation in household consumption. Specifically, we look at TRA86, which is bound to boost consumption, and OBRA93, which increased top marginal tax rates thus potentially depressing consumption. The use of both a positive and negative shock allows us to test the following hypotheses regarding the social preference parameter: (i) Is θ a short- or long-run parameter? (ii) are 'keeping-up effects' downward rigid, in the sense that households – having once increased consumption for the sake of pecuniary emulation – are reluctant to cut back on expenditures when faced with a negative income shock?

We obtain data on household income, consumption, and demographic characteristics from the family-level survey extracts of the *Consumer Expenditure Survey* (*CEX*). Our sample spans 1983:Q1-1989:Q4 to test the effect of TRA86, and 1990:Q1-1997:Q4 to test the effect of OBRA93. Household consumption is calculated as the sum of household expenditures on food, clothing, rent and utilities, motor vehicles and parts, furniture, housing intermediate goods, life insurance premiums, non-food items such as tobacco and alcohol, transportation costs (such as gasoline), education expenditures, recreation (including items such as books or gambling expenses), donations to charity, and out of pocket medical expenses (including cost of health insurance). While providing a fairly comprehensive definition of total household consumption, it is nonetheless well known that the CEX tends to underestimate total consumption, especially of high income households (Cooper, 2010; Sablehaus, 2010; Caroll et al., 2015). This means that—to the extent that the consumption responses of high income households to changes in taxation are understated in the data—our estimate of the social preference parameter can be thought of as providing a lower-bound for the true value.

The final sample includes only households who satisfy three criteria. First, households must have met the 'complete income reporter' requirement by the Bureau of Labor Statistics. Second, households must have completed all four consecutive quarterly interviews. Finally, student households are dropped from the sample. The sample is then re-weighted to adjust for attrition. For each household, all four quarterly interviews are merged into an annual record. Thus, the data take the form of repeated cross-section, where each household reports an annualized quarterly value for the variable of interest. All income and expenditure variables are converted to constant 1999 dollars. Appendix D presents summary statistics and further discussion of the data.

In order to obtain an estimate of the social preference parameter, we use the following two-stage procedure. In the first stage, we exploit the exogenous variation in consumption caused by federal tax reform using a difference-in-differences approach:

$$c_{it} = \alpha_0 + \alpha_1 TaxTreat_i + \alpha_2 After_t + \alpha_3 (TaxTreat \times After)_{it} + \alpha_4 y_{it} + v_{it} \quad \text{(1S)}$$

Where $TaxTreat_i$ takes a value of 1 if the observed household is in the group treated by the tax reform and 0 otherwise. Specifically, we focus on households who fall in the top tax bracket prior to the reform. The main effect of both TRA86 and OBRA93 is to alter the marginal tax rates facing top income earners, suggesting it is these households whose consumption patterns are most likely to be altered post-reform. $After_t$ takes a value of 1 after the reform and 0 otherwise, y_{it} reports household income, and v_{it} is an idiosyncratic shock.

Using the parameter estimates from (1S), we obtain predicted values for household consumption, $\hat{c}_{it}$. Using these values, we calculated predicted values for mean within-cohort consumption, partitioning households into cohorts defined as deciles of the income distribution:

$$\hat{\hat{c}}_{jt} = \frac{1}{n_j} \sum_{i \in j} \hat{c}_{it} \tag{24}$$

where n_j is the number of households in cohort j. In the second stage, $\hat{\hat{c}}_{jt}$ is then used as an instrument in (2S) to obtain an estimate of θ. We apply this method first for TRA86 and then for OBRA93.

4.1. Tax Reform Act of 1986

For the difference-in-differences approach to serve as a valid strategy for obtaining estimates of $\hat{\hat{c}}_{jt}$, the consumption-path of households facing the top marginal tax rate must satisfy the parallel trend assumption with the control group in the pre-treatment period (in this case, all households outside the top tax bracket). Unfortunately, in the case of TRA86 it is not clear when the full effect of the treatment becomes salient to households. TRA86 is passed in 1986:Q4, meaning households face new rates beginning with tax season 1987. However, beginning in 1988 a 'tax bubble' is phased in, altering the effective rates paid by households. As a result of the tax bubble, households earning between approximately $75,000 and $170,000 paid a 33% top marginal rate, while households earning even more paid only a 28% rate. To try and capture the effect of the tax bubble, we estimate values of θ for three different post-treatment periods: after 1986:Q4, 1987:Q4, and 1988:Q1. This ensures that we control for any potential tax salience effects on household consumption. Figure 2 presents trend plots with regression discontinuities at each of the above dates.

[Figure 2 about here.]

The trends appear most parallel through 1987:Q4 or 1988:Q1. This is not surprising, given that the full effect of the tax reform may not have been salient to households until the tax bubble was fully phased in. Nonetheless, we make use of all three breaks in our estimation of θ as a robustness check. Table 1 presents results from the first-stage difference-in-differences estimation.

[Table 1 about here.]

The results indicate that for two of the three post-treatment period definitions TRA86 had a significant effect on the consumption of households falling in the top marginal tax bracket. For the post-1987:Q4 period, the treatment effect is economically large, but slightly statistically insignificant (t-stat = 1.51). The coefficient on y_{it} is constant across all specifications, indicating a relatively stable relationship between income and consumption for alternative definitions of the pre- and post-treatment period.

Using the estimated coefficients from (1S), we obtain predicted values for mean cohort consumption, $\hat{\bar{c}}_{ij}$ according to (24). These values are then used to estimate (2S). The standard errors of this second stage regression can then be computed via bootstrapping. To capture household- and time-specific heterogeneity, we use time-fixed effects and a vector of controls for household demographic characteristics including age, race, gender, education level, and marital status of the household head. As a measure of household savings we include the sum of the total amounts in the household's checking and saving accounts. Although an imperfect measure, some control for household saving is important – both because of the role saving plays along balanced growth paths, and for what the relationship between savings and consumption tells us about the rate of time preference in the economy. The results of the estimation are presented in Table 2.

[Table 2 about here.]

Contingent on identification in the first-stage regression, $\hat{\bar{c}}_{jt}$ is a valid instrument for $\bar{c}_{jt}$. Given significance of the treatment effect co-efficient and the size of the first-stage F statistic, this appears to be the case. Thus, given that the $\hat{\bar{c}}_{jt}$ satisfies the desired IV properties, what can be said about the estimated θ? Our estimate is both statistically ($p < 0.01$) and economically significant across all post-treatment period definitions. The coefficient on $\bar{c}_{ij}$ is constant across all definitions of the post-treatment period. This is not terribly surprising given that sample mean estimates of the predicted value of $\hat{\bar{c}}_{ij}$ are relatively constant from within-cohort across definitions of the post-treatment period. The implication is that, on the one hand, the instrument being used in each case is essentially the same, with only the slightest variation resulting from different definitions of the post-treatment period. On the other hand, salience effects on social consumption behavior due to the tax bubble are small compared to the overall effect of the change in the tax schedule.

The estimated coefficient on $\hat{c}_{ij}$ tells that, if average consumption in income cohort j increases by \$10, then household i in cohort j will increase its consumption by approximately \$3.30. Intuitively, our estimate of θ answers the question: if the people on my rung of the social ladder consume an extra dollar, does that impact my consumption? It is thus an estimate of the marginal propensity to consume out of average cohort consumption. The answer to the above question appears to be yes, in a non-trivial manner. In the theoretical model, the representative household is also the average household along a BGP. Thus, combined with the coefficient on y_{it} —the marginal propensity to consume out of income— the societal marginal propensity to consume is approximately 0.55 in the benchmark AK model.[4]

4.1.1. Robustness Checks — TRA86

One possible confounding factor in the above estimation is the inclusion of households reporting negative or zero incomes in the income cohort at the left tail of the distribution. Such households may report negative income due to capital losses on owned assets, report no income by mistake, or receive zero income despite large holdings of wealth. If any of these scenarios occur, then inclusion of these households in the same cohort as those households truly at the bottom of the socioeconomic ladder may have confounding effects, as it is not clear that these two groups are in fact best responding to one another. To test this, we run our two-stage estimator on a restricted sub-sample that excludes households reporting incomes less than or equal to zero. We make use the 1986:Q4 treatment period definition for this test. Table 3 presents the results.

[Table 3 about here.]

The results from the first-stage difference-in-differences estimation are similar to those presented previously. The estimate of θ, while still significant, is somewhat smaller than previous estimations, indicating that the relationship between households reporting an income less than or equal to 0 and the rest of those households in or near the left-tail of the income distribution influences the magnitude of θ.

Another potential objection is that the strength of social preferences in consumption may vary across groups within the income distribution. Because it presents an economy-wide average, θ captures the magnitude of the average effect of social preferences on consumption across income cohorts. It is then possible that this effect is comparatively stronger for

[4] Perhaps suggestively, from these coefficients one can calculate a textbook autonomous spending multiplier $1/(1 - MPC)$ around 2.2.

some cohorts than others. In particular, the ‘keeping up’ story is typically told in reference to the middle class, but whether or not this holds in reality requires verification.

To test for differences across income cohorts, we make use of Amemiya’s (1986) two-stage least absolute deviation (2SLAD) estimator for quantile regression in the presence of endogenous variables. We use quantile regression to obtain estimates of the magnitude of θ for different percentiles. Table 4 presents the results. We use bootstrapping to ensure our standard errors are appropriate given the estimation procedure. The use of quantile regression means our estimates will capture the effect of a change in $\overline{c}_{jt}$ on different percentiles of the response variable (that is, individual consumption). In this case, the distribution we are stratifying across is technically the distribution for consumption, rather than income. However, we see this as an advantage, as it avoids any potential bias that could result from grouping individuals who either under reported or failed to report income—but whom report large values for consumption, possibly due to wealth—with those who “truly” belong in the left tail of the consumption-income distribution.

[Table 4 about here.]

The results indicate that our intuition regarding the variation in social preferences among different socioeconomic groups was partially correct. Note that due to the tendency of the CEX to understate the consumption of high income households, the estimates of θ for the upper income quintiles will tend to be biased downward from their true value. However, it is possible that the consumption data in the CEX reasonably approximate the consumption of ‘keeping up’ goods, such as automobiles, housing durable goods, and expenditures on food and clothing, such that the size of this bias is likely small. Our estimate of the extent of social preferences is monotonically decreasing throughout the consumption distribution, taking on negative —albeit not statistically significant— values for those individuals at the very top of the income ladder. We interpret this as suggesting that magnitude of social preferences in consumption is essentially zero as one approaches the top of the distribution. The strength of θ in the lower tail of the distribution accords with some notion of a ‘pecuniary standard of living’ which Veblen (1899) argued set a relative standard of consumption for which anyone who fell below it would ceaselessly strive to achieve.

Lastly, when a treatment effects approach is used to identify the causal relation of interest, the model can be subjected to a ‘placebo’ test. To conduct a placebo test, an earlier period —where it is known that the actual treatment (in this case, TRA86) has not yet been administered— is assigned as the treatment date to the same treatment group. If the effect of the policy change is identified, then there should be no significant behavioral responses to the placebo. We adopt this approach by assigning an alternative date of 1984:Q2, which we know to be before the change in the tax code. We restrict the sample to the entire period before 1986:Q3 to avoid confounding effects from the actual treatment. Table (5) presents the results.

[Table 5 about here.]

The coefficient on $(After \times TaxTreat)_{it}$ is highly statistically insignificant and economically small in comparison to the actual treatment effect. The results of the placebo test thus offer further support for the appropriateness of our identification strategy.

A further question regards the relationship between the direction of the exploited policy change and θ. If there exists some third factor (e.g. consumer optimism) that is correlated with both the direction of the policy change and household consumption, then our estimates of θ may be biased and incorrect. Given the expansionary effect of the TRA86 tax cuts, this is not wholly implausible. Furthermore, it may be that social preferences in consumption vary in time, in which case a single estimate of θ will not prove a valuable guide to policy. Whether or not θ is constant over time is at least as important as the magnitude of θ at a given point in time. Finally, if the magnitude of social preferences in consumption is less for negative shocks than for positive shocks – i.e. household consumption is socially downward rigid, in the sense that households, once reaching a certain socially sanctioned level of consumption, are reluctant to cut back – this will have important policy implications. To address these questions, we turn to the Omnibus Budget Reconciliation Act of 1993.

4.2. Omnibus Budget Reconciliation Act of 1993

Passed in August of 1993, OBRA93 increased marginal tax rates for top income earners. In particular, the top bracket was split in two, and the marginal rates were raised from 0.31 to 0.36 and 0.396, respectively. While these changes are small relative to the magnitude of the changes that occurred under TRA86, they are nonetheless salient enough to provide a negative counterpart to the positive consumption shock induced by the 1986 reform. We again make use of the Consumer Expenditure Survey family-level extracts, with data spanning 1990:Q1-1997:Q4.

[Figure 3 about here.]

Figure 3 presents a simple parallel trends plot for the treatment and control groups for OBRA93. There appears to be a slight downward shift in the consumption trend of the treated group after 1993:Q3, which however is relatively minor. Unlike the case of TRA86, regression discontinuity analysis proved inconclusive, as the results were highly sensitive to the time window selected around the reform. Additionally, the effect of OBRA93 on consumption is smaller relative to the overall trend in consumption than in the case of TRA86. The results of the regression discontinuity analysis are included in the appendix as a sensitivity check.

Given the parallel trends assumption is satisfied, we again use a difference-in-differences approach to obtain an estimate of $\bar{\hat{c}}_{jt}$ via (1S) to use as an instrument in (2S). Table 6 presents the results from the difference-in-differences regression.

[Table 6 about here.]

The treatment effect has the expected negative sign and is economically and statistically significant. The treatment, that is the change in the tax schedule resulting from OBRA93, resulted in an estimated decline in consumption of the treated group between $-\$2,437$ and $-\$8,633$ in constant 1999 dollars, given the standard error of the regression coefficient.

Using the coefficients estimated from applying (1S) to OBRA93, we again obtain predicted values for mean cohort consumption, $\bar{\hat{c}}_{jt}$ via (24). Finally, these values are again used as an instrument to estimate (2S). The results are displayed in Table 7.

[Table 7 about here.]

The estimate of θ is both statistically and economically significant, while the magnitude is slightly smaller than in the case of TRA86. This could be for several reasons. First, it is possible that the size of the change in the marginal tax rate elicited a smaller behavioral response relative to TRA86. Second, it is possible that households are in fact downward rigid in their socially-informed consumption behavior and that individual consumption is less sensitive to movements in average consumption when that movement is negative. This would support the idea that —having increased consumption for reasons of pecuniary emulation— households are reluctant to cut back for fear of failing to 'keep up.' Finally, given the standard error of the estimates on θ for both TRA86 and OBRA93, it would not be unreasonable to conclude that in each case θ falls within the same plausible range of implied values. A test of equality that $\theta_{obra93} = \theta_{tra86}$ fails to reject the null.[5]. This result points to the conclusion that θ is relatively constant over time, and can be treated as a long-run parameter of the economy.

4.2.1. Robustness Checks — OBRA93

The same critiques that apply to our estimation in the case of TRA86 apply to OBRA93. The inclusion of zero income reporters may bias results if those households are not true zeros in a socioeconomic sense. We again check for robustness by running our estimation on a sub-sample for the OBRA93 period that excludes zero income reporting households. Table 8 reports the results.

[Table 8 about here.]

As before, the magnitude of θ is somewhat attenuated as compared to when the zero income reporters are included, although the reduction in the case of OBRA93 is smaller than for TRA86 (likely due to the smaller number of zero income reporters —40 in the OBRA93 sample, as compared with 78 in the previous case).

The second critique regards the possibility of differential magnitudes of θ across cohorts of the income distribution. To test for these differences we again estimate θ for different parts of the income distribution using a quantile regression as before. A priori, we again expect the 'keeping up' effect to be strongest in the lower and middle region of the consumption distribution as those individuals near the middle-class strive to achieve the measure of social status that comes with reaching a particular level of consumption. It is not wholly unreasonable that certain households feel, as Veblen (1899) writes, "[I]n order to stand well in the eyes of the community, it is necessary to come up to a certain, somewhat indefinite, conventional standard of wealth" (25). The results are presented in Table 9.

[5] Because the models are non-nested the test statistic is, however, imperfectly estimated. The estimated statistic is $z = \frac{\theta_{tra86} - \theta_{obra93}}{\sqrt{(SE\theta_{tra86})^2 + (SE\theta_{obra93})^2}}$, which doesn't take into account possible co-variance between the parameter estimates. However, the magnitude of the test statistic ($z \approx 0.71$) is small enough to belay concern.

[Table 9 about here.]

The results conform to those from the TRA86 case. The effect of θ is decreasing in the percentile of the consumption distribution, with the largest effects of $\overline{c}_{jt}$ on individual consumption occurring below the 40th percentile.

Finally, we conduct a placebo test for OBRA93. We assign a treatment date of 1992:Q2 and restrict the sample to the period prior to the implementation of OBRA93. Table (10) presents the results. Once again, we observe no significant effect of the placebo on consumption.

[Table 10 about here.]

5. Conclusion

Social preferences matter. From our empirical results, it appears that on average households are heavily influenced by peer consumption standards in their spending decisions, and that these effects are consistent with a long-run interpretation of preferences being other-regarding. If endogenous growth models provide an adequate representation of the growth process, as the evidence found by Bernanke and Gürkaynak (2001) seem to point out, then 'keeping up' behavior has non trivial effects in the long run in that it permanently reduces the growth rate of the economy. Equally interesting is the result on semi-endogenous growth models: 'keeping up' behavior unambiguously increases the long-run number of R&D workers in the economy. Because of diminishing returns to existing ideas, many of these additional scientists might end up engaged in duplication activity instead of contributing to long-run growth.

Another result that stands out from our empirical analysis is that the extent of social preferences is not symmetric across all quintiles of the population, but tends to vanish as incomes increase. The implication is that even if households have identical time preference rates across the entire income distribution, those at the top of the income distribution will see their share of wealth increase over time, because they are not as preoccupied about consuming to 'keep up' with their peers as lower-income households are. Thus, our empirical analysis suggests a channel through which income inequality and wealth inequality are related.

References

Acemoglu, D., 2009. *Introduction to Modern Economic Growth*, Princeton.

Agarwal, S., and Qian, W. 2014. "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore." *American Economic Review*, 104(12): 4205-4230.

Aghion, P., and Howitt, P., 1992. A Model of Growth through Creative Destruction. *Econometrica*, 60: 323-351.

Alvarez-Cuadrado, F., and Van Long, N. (2011). The Relative Income Hypothesis. *Journal of Economic Dynamics and Control*, 35: 1489171501.

Amemiya, T., 1982. Two-Stage Least Absolute Deviation Estimators. *Econometrica*, 50(3): 689-711.

Barro, R. J. (1990) Government spending in a simple model of endogenous growth, *Quarterly Journal of Economics* 98 (5): S103-125.

Bernanke, B. S., and Gürkaynak, R. S. (2001). Is Growth Exogenous? Taking Mankiw, Romer, and Weil Seriously. In Bernanke, B. S., and Rogoff, K., eds: *NBER Macroeconomics Annual 2001*, Vol. 16. Cambridge, MA: MIT Press.

Carroll, C.D., Overland, J., and Weil, D.N. (1997). Comparison Utility in a Growth Model. *Journal of Economic Growth*, 2(4): 339-367.

Caroll, C.D., Crossley, T.F., and Sabelhaus, J., eds. (2015). *Improving the Measurement of Consumer Expenditures*, 74: University of Chicago Press.

Cooper, D. (2010). "Imputing Household Spending in the Panel Stud of Income Dynamics: A Comparison of Approaches." *Federal Reserve Bank of Boston*, Working Paper 10-12.

Duesenberry, J. S., 1949. *Income, Saving and the Theory of Consumer Behaviour*, Harvard University Press.

Dynan, K. E., and Ravina, E., 2007. Increasing Income Inequality, External Habits, and Self-Reported Happiness. *American Economic Review* 97(2): 226-231.

Grossman, M. and Helpman, E. (1991) Quality ladders in the theory of growth, *Review of Economic Studies* 58 (1): 43-61.

Jones, C. (1995) R&D-based models of economic growth. *Journal of Political Economy* 103 (4): 759-784.

Liu, W.-F., and Turnovsky, S.J.,2005. Consumption Externalities, Production Externalities, and Long-run Macroeconomic Efficiency. *Journal of Public Economics* 89(5-6): 1097-1129.

Ljungqvist, L.,and Uhlig, H.,2000. Tax Policy and Aggregate Demand Management under Catching up with the Joneses. *American Economic Review* 90, 356-366.

Lucas, R.J., 1988. On the Mechanics of Economic Development, *Journal of Monetary Economics*,

Misra, K., and Surico, P. 2014. "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *AEJ: Macroeconomics*, 6(4): 8-106.

Ravina, E., 2007. Habit Formation and Keeping Up with the Joneses: Evidence from Micro Data. Available at SSRN: http://dx.doi.org/10.2139/ssrn.928248.

Romer, P. M., 1986. Increasing Returns and Long-Run Growth. *Journal of Political Economy* 94: 1002-37.

Romer, P. M., 1990. Endogenous technological change. *Journal of Political Economy* 98 (5) S71-S102.

Sablehaus, J. (2010). The Joint Distribution of Consumption and Income, working paper, University of Maryland.

Stock, J., and Staiger, D. 1997. Instrumental Variables Regression with Weak Instruments. *Econometrica*, 65(3): 557-586.

Stock, J., Wright, J., and Yogo, M. A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments. *Journal of Business and Economic Statistics*, 20(4): 518-529.

Stock, J., and Yogo, M. 2005. Testing for Weak Instruments in Linear IV Regression, in Andrews, D. & Stock, J., eds. *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge: Cambridge University Press.

Turnovsky, S., Monteiro, G., and Alvarez-Cuadrado, F. (2004). "Habit Formation, Catching up with the Joneses, and Economic Growth." *Journal of Economic Growth*, 9: 47-80.

Veblen, T., 1899 [2007]. *The Theory of the Leisure Class*, Oxford: Oxford University Press.

A CRRA Preferences

The best-response function modifies to:

$$c(\bar{c},\theta) = \frac{[(\gamma-1)A+\rho]k}{2\gamma} + \frac{\left\{\gamma^{-2}[(\gamma-1)A+\rho]^2k^2 + 4\theta\bar{c}k[A-\rho\right\}^{1/2}}{2} \tag{25}$$

which has intercept equal to $\gamma^{-1}\left[(\gamma-1)A+\rho\right]k$ and is increasing and concave in average consumption. The equilibrium and efficient growth rates are respectively

$$g = \frac{1-\theta}{\gamma}[A-\rho] < g^* = \gamma^{-1}[A-\rho], \quad \forall\theta \in (0,1); \tag{26}$$

while the equilibrium and efficient consumption are, respectively

$$c = \left[\frac{\gamma-1+\theta}{\gamma}A + \frac{1-\theta}{\gamma}\rho\right]k > c^* = \left[\frac{\gamma-1}{\gamma}A + \frac{\rho}{\gamma}\right]k, \quad \forall\theta \in (0,1). \tag{27}$$

B Speed of Convergence in the NGM with Other-Regarding Preferences

Linearization around the steady state results in the following Jacobian matrix:

$$J_{ss} = \begin{pmatrix} 0 & -(1-\theta)\xi \\ -1 & \rho \end{pmatrix},$$

where $\xi \equiv (1-\alpha)(\alpha A)^{1/(\alpha-1)}\rho^{\frac{\alpha-2}{\alpha-1}} > 0$. The linearized system, thus, has the typical saddle-path stability property. The eigenvalues of J_{ss} are

$$\lambda_{1,2} = \rho \pm \frac{\left[\rho^2 + 4(1-\theta)\xi\right]^{1/2}}{2},$$

of which the negative root, say λ_1, is the speed of convergence to the steady state. We have that $\partial\lambda_1/\partial\theta > 0$ which, since λ_1 is negative, indicates that social preferences determine convergence to the steady state at a slower pace than in the selfish case.

C Additional Sensitivity Analysis

C1. OBRA93 - Regression Discontinuities

[Figure 4 about here.]

Figures 4a and 4b present regression discontinuity plots for OBRA93. The sensitivity of the analysis to the choice of time-window around the treatment is evident. The former presents a regression discontinuity for a reduced time-window as compared to the whole sample. However, we believe that the reduced-window is likely a more accurate depiction of the impact of the treatment. It is unlikely that the consumption effects of tax reform persist longer than three years, and extending the window any further back in time may reflect earlier shocks and/or shifts in the trend for consumption. Thus, while not as conclusive as the RD plots for TRA86, we believe that Figure 4a provides a reliable representation of the impact of OBRA93.

C2. Lagged Consumption

Other studies of consumption peer effects have used some measure of lagged average consumption as the benchmark for an individual consumer. This section reports the results from a similar test using the main specification in our study, where instead of predicted contemporaneous average consumption we use a one-period lagged value of predicted average consumption as the instrumental variable. Table (11) presents the results for both TRA86 and OBRA93.

[Table 11 about here.]

D Data

Tables (12) and (13) present summary statistics for the TRA86 and OBRA93 samples, respectively.

[Table 12 about here.]

[Table 13 about here.]

Tables

Table 1: Difference-in-Differences Estimation — Tax Reform Act of 1986

	After 1986:Q4	After 1987:Q4	After 1988:Q1
y_{it}	0.424***	0.424***	0.424***
	(0.00648)	(0.00647)	(0.00645)
$After_{it}$	90.93	513.7**	574.0**
	(302.4)	(316.5)	(320.5)
$TaxTreat_{it}$	-33330.07***	-31132.5***	-33219.9***
	(5427.1)	(5102.5)	(4978.2)
$(After \times TaxTreat)_{it}$	14481.1*	13466.3	19542.9**
	(8340.8)	(8926.4)	(8527.2)
N	20242	20242	20242
R^2	0.483	0.484	0.484
F Statistic	1196.7	1156	1147.4

Standard errors in parentheses, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 2: Estimates of θ — TRA86

	After 1986:Q4	After 1987:Q4	After 1988:Q1
$\bar{c}_{jt}$	0.327***	0.327***	0.327***
	(0.0555)	(0.0578)	(0.0538)
y_{it}	0.219***	0.220***	0.219***
	(0.0250)	(0.0246)	(0.0234)
S_{it}	0.0315***	0.0315***	0.0315***
	(0.00718)	(0.00714)	(0.00664)
Controls	Y	Y	Y
Time FE	Y	Y	Y
N	20242	20242	20242
R^2	0.523	0.523	0.523

Bootstrapped standard errors in parentheses, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 3: Restricted Sample — No Zero Income Reporters — TRA86

	Diff-in-Diff	IV
y_{it}	0.432***	0.241***
	(0.00603)	(0.0235)
$After_t$	216.3	—
	(280.7)	—
$TaxTreat_i$	-34724.7***	—
	(5428.9)	—
$(After \times TaxTreat)_{it}$	14495.6*	—
	(8376.8)	—
$\overline{c}_{jt}$	—	0.287***
	—	(0.0525)
S_{it}	—	0.0262***
	—	(0.00598)
N	20164	20164
R^2	0.498	0.531
F Statistic	1471.9	—

Standard errors in parentheses, clustered at the cohort-quarter level. IV errors estimated via bootstrapping. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 4: θ, Stratified by percentile — TRA86

	5th	20th	40th	60th	80th	95th
$\bar{c}_{jt}$	0.405***	0.326***	0.306***	0.248***	0.154**	-0.0596
	(0.0348)	(0.0450)	(0.0429)	(0.0482)	(0.0643)	(0.0986)
Controls	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y
N	20164	20164	20164	20164	20164	20164
Pseudo-R^2	0.29	0.35	0.37	0.38	0.37	0.36

Bootstrapped standard errors in parentheses. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 5: Placebo Test — TRA86

	After 1984:Q2
y_{it}	0.419***
	(0.0109)
$After_{it}$	29.91
	(474.3)
$TaxTreat_{it}$	-35613.0***
	(3368.6)
$(After \times TaxTreat)_{it}$	7886.8
	(10860.5)
N	10056
R^2	0.476
F Statistic	448.18

Standard errors in parentheses, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 6: Difference-in-Differences Estimation — OBRA93

	After 1993:Q3
y_{it}	0.410***
	(0.00746)
$After_t$	801.7***
	(247.0)
$TaxTreat_i$	-6038.9**
	(2840.4)
$(After \times TaxTreat)_{it}$	-5534.6*
	(3097.7)
N	22686
R^2	0.488
F Statistic	1041.2

Standard errors in parentheses, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 7: Estimates of θ — OBRA93

$\overline{c}_{jt}$	0.277***
	(0.0608)
y_{it}	0.241***
	(0.0246)
S_{it}	0.0411***
	(0.00629)
Controls	Y
Time FE	Y
N	22686
R^2	0.523

Bootstrapped errors in parentheses, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 8: Restricted Sample — No Zero Income Reporters — OBRA93

	Diff-in-Diff	IV
y_{it}	0.414***	0.248***
	(0.00752)	(0.0276)
$After_t$	756.9***	—
	(245.2)	—
$TaxTreat_i$	-6533.5**	—
	(2848.8)	—
$(After \times TaxTreat_{it}$	-5496.1*	—
	(3100.5)	—
$\bar{c}_{jt}$	—	0.262***
	—	(0.0667)
S_{it}	—	0.0409***
	—	(0.00686)
Controls	N	Y
Time FE	N	Y
N	22646	22646
R^2	0.493	0.526
F Statistic	1038.6	—

Standard errors in parentheses, clustered at the cohort-quarter level. IV errors estimated via bootstrapping. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 9: θ, Stratified by percentile — OBRA93

	5th	20th	40th	60th	80th	95th
$\bar{c}_{jt}$	0.334***	0.317***	0.266***	0.250***	0.117**	0.0395
	(0.0414)	(0.0389)	(0.0460)	(0.0473)	(0.0471)	(0.175)
Controls	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y
N	22646	22646	22646	22646	22646	22646
Pseudo R^2	0.27	0.33	0.36	0.37	0.37	0.37

Bootstrapped standard errors in parentheses. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 10: Placebo Test — OBRA93

	After 1992:Q2
y_{it}	0.410***
	(0.00747)
$After_{it}$	640.2**
	(265.0)
$TaxTreat_{it}$	-7438.4*
	(4143.3)
$(After \times TaxTreat)_{it}$	-2804.5
	(4158.0)
N	22686
R^2	0.488
F Statistic	951.66

Standard errors in parentheses, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table 11: Estimates of θ - Lagged Consumption - TRA86, OBRA93

	TRA86	OBRA 93
$\bar{c}_{j,t-1}$	0.29***	0.26***
	(0.05)	(0.07)

Notes: Bootstrapped standard errors in parenthesis, clustered at the cohort-quarter level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Table 12: Sample Means — 1983:Q1 - 1989:Q4 (TRA86)

	Mean	Std. Dev
y_{it}	\$43,061.22	35,225.23
c_{it}	\$32,294.34	20,595.52
S_{it}	\$7,698.62	19,112.53
Black	0.11	0.31
Native American	0.01	0.11
Asian	0.01	0.10
White	0.87	0.34
$<$ High School	0.36	0.48
High School	0.27	0.44
Some College	0.16	0.37
College	0.11	0.31
Grad School	0.09	0.29
Female	0.33	0.47
Dual	0.005	0.07
Age	47	17
N	20242	

Education levels refer to the highest level completed by the household head. Female and Dual refer to whether or not the household is a female- or dual-headed household, respectively. Race variables refer to the race of the household head. Dollar values in constant 1999 $.

Table 13: Sample Means — 1990:Q1 - 1997:Q4 (OBRA93)

	Mean	Std. Dev
y_{it}	$43,084.03	20,729.14
c_{it}	$31,692.02	36,797.07
S_{it}	$9,064.9	26,510.78
Black	0.11	0.31
Native American	0.005	0.08
Asian	0.02	0.15
White	0.86	0.35
$<$ High School	0.20	0.40
High School	0.32	0.46
Some College	0.23	0.42
College	0.15	0.35
Grad School	0.10	0.30
Female	0.38	0.48
Dual	0.006	0.08
Age	48	18
N	22686	

Education levels refer to the highest level completed by the household head. Female and Dual refer to whether or not the household is a female- or dual-headed household, respectively. Race variables refer to the race of the household head. Dollar values in constant 1999 $.

Figures

Figure 1: Equilibrium *vs.* efficient consumption.

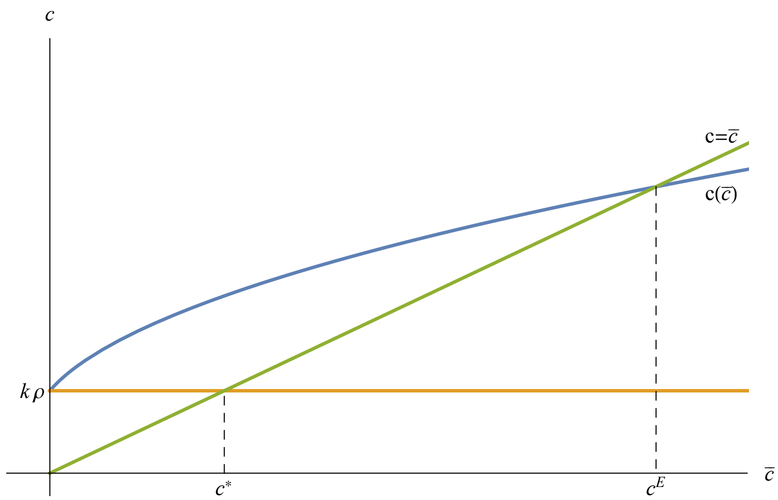

Figure 2: TRA86 — RD Plots

(a) Regression Discontinuity - 1986:Q4

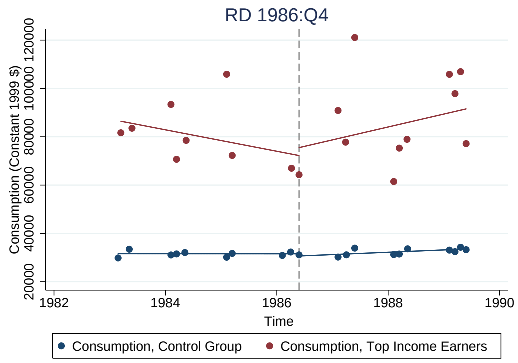

(b) Regression Discontinuity - 1987: Q4

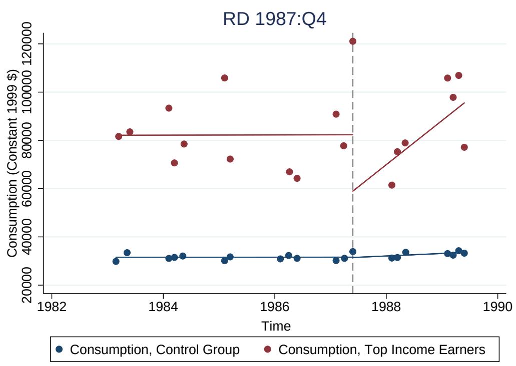

(c) Regression Discontinuity - 1988:Q1

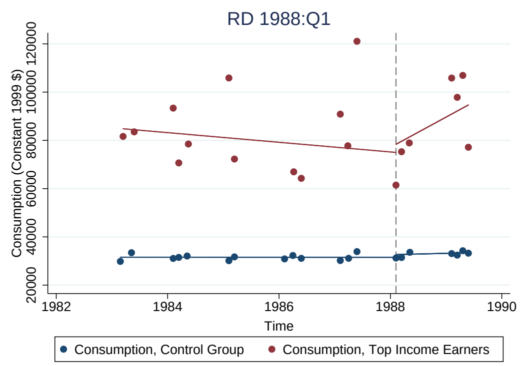

Figure 3: OBRA93 — Parallel Trends

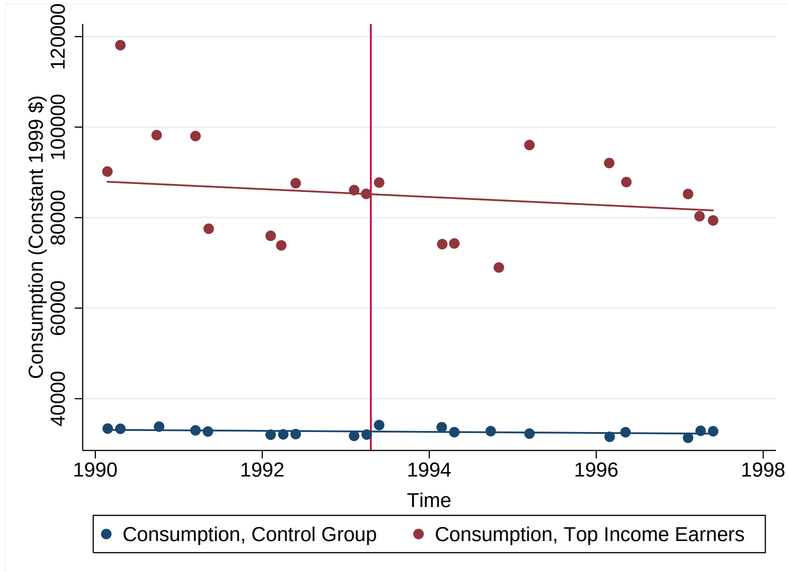

Figure 4: OBRA93 - Regression Discontinuities

(a) Regression Discontinuity - 1991-1996

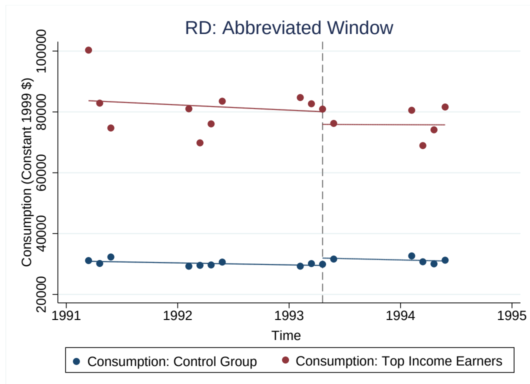

(b) Regression Discontinuity - Full OBRA93 Sample